REMARKS

In view of the following discussion, the Applicants submit that none of the claims now pending in the application is directed to non-statutory subject matter under the provisions of 35 U.S.C. §101 or is anticipated under the provisions of 35 U.S.C. §102. Thus, the Applicants believe that all of these claims are in allowable form.

I. REJECTION OF CLAIMS 1-2, 7-8, AND 13-14 UNDER 35 U.S.C. § 101

A. Claims 1-2

Claims 1-2 stand rejected under 35 U.S.C. § 101 as being allegedly directed to non-statutory subject matter. In response, the Applicants have amended independent claim 1 in order to more clearly recite aspects of the invention.

In particular, independent claim 1 has been amended to recite that "at least one of: the receiving, the identifying, the updating a threshold similarity, the updating a similarity expectation, the comparing, the associating, or the defining is performed by a processor" (emphasis added). As such, independent claim 1 clearly recites a method that is tied to a particular machine or apparatus (i.e., a processor) that performs the recited steps. The "Interim Examination Instructions For Evaluating Subject Matter Eligibility Under 35 U.S.C. § 101," effective August 24, 2009, provide that a process comprising "an act, or a series of acts or steps that are tied to a particular machine or apparatus" constitutes patent-eligible subject matter (See, Interim Examination Guidelines, Page 1, "Subject Matter Eligibility"). As such, the Applicants respectfully submit that independent claim 1 is directed to subject matter that is statutory within the meaning of 35 U.S.C. § 101.

Claim 2 depends from independent claim 1 and recites at least all of the features recited in independent claim 1. As such, and at least for the reasons stated above with respect to independent claim 1, the Applicants respectfully submit that claim 2 is also directed to subject matter that is statutory within the meaning of 35 U.S.C. § 101.

In light of the above, the Applicants respectfully submit that claims 1-2 fully satisfy the requirements of 35 U.S.C. §101. Accordingly, the Applicants respectfully request that the rejection of claims 1-2 under 35 U.S.C. §101 be withdrawn.

B. Claims 7-8

Claims 7-8 stand rejected under 35 U.S.C. § 101 as being allegedly directed to non-statutory subject matter. In response, the Applicants have amended claims 7-8 in order to more clearly recite aspects of the invention.

In particular, claims 7 and 8 have been amended to explicitly recite a "computer readable storage medium" (emphasis added) that contains the executable program. A computer readable storage medium (e.g., random access memory, magnetic drive, optical drive, diskette, or the like) is clearly statutory subject matter within the meaning of 35 U.S.C. § 101. Specifically, "[i]n this context, 'functional descriptive material' consists of data structures and computer programs which impart functionality when employed as a computer component." (MPEP 2106.01) "When functional descriptive material is recorded on some computer-readable medium, it becomes structurally and functionally interrelated to the medium and will be statutory in most cases since use of technology permits the function of the descriptive material to be realized." (MPEP 2106.01)

Since the claimed executable program is contained on a computer readable storage medium, the executable program is "structurally and functionally interrelated" to the computer readable storage medium, and, as such, is statutory in accordance with MPEP 2106.01.

Moreover, the Board of Patent Appeals and Interferences recently noted in Ex parte Bo Li, 2008-1213, that "Beauregard" claims are viewed as product claims, and are therefore statutory under 35 U.S.C. §101. See Bo Li at Page 9 ("It has been the practice for a number of years that a 'Beauregard Claim' ... be considered statutory at the USPTO as a product claim. (MPEP 2105.01, I). Though not finally adjudicated, this practice is not inconsistent with In re Nuijten. (Ibid)"). The Board further noted that a combination of software components embodied upon a computer readable medium "has been found statutory under the teachings of In re Lowry, 32 F.3d 1579 (Fed. Cir. 1994)." (Bo Li at Page 9). The Applicants respectfully submit that the holding of the Board in Ex parte Bo Li therefore supports a finding that the subject matter embodied in claims 7-8

is statutory.

In light of the above, the Applicants respectfully submit that claims 7-8 fully satisfy the requirements of 35 U.S.C. §101. Accordingly, the Applicants respectfully request that the rejection of claims 7-8 under 35 U.S.C. §101 be withdrawn.

C. Claims 13-14

Claims 13-14 stand rejected under 35 U.S.C. § 101 as being allegedly directed to non-statutory subject matter. The Applicants respectfully traverse the rejection.

In particular, the Applicants respectfully disagree with the Examiner's allegation that the invention recited in claims 13-14 "recite[ ] only abstractions that are neither 'things' nor 'acts'" (Office Action, Page 5). The Applicants submit that claims 13-14 and 13-14 are clearly directed towards an apparatus (*i.e.*, an intrusion detection system) and are statutory *per se* under 35 U.S.C. § 101. In particular, the Applicants submit that the necessary physical articles are embodied at least in the recited "means for receiving," "means for identifying," "means for updating a threshold similarity," "means for updating a similarity expectation," "means for comparing," and "means for associating" in claim 13. 35 U.S.C. 112, sixth paragraph allows that "[a]n element in a claim for a combination may be expressed as a <u>means or step for performing a specified function</u> <u>without the recital of structure, material, or acts in support thereof</u>, and <u>such claim shall</u> <u>be construed to cover the corresponding structure, material, or acts described in the</u> <u>specification and equivalents thereof</u>" (emphasis added). Accordingly, a patentee may generically define a structure for performing a particular function through the use of a means expression, provided that the specified structures corresponding to the means are disclosed in the patent specification. <u>Warner-Jenkinson Co., Inc. v. Hilton Davis</u> <u>Chemical Co.</u>, 520 U.S. 17, 28 (1997).

As claim 13 recites means for performing specified functions, the Applicants respectfully submit that claim 13 is a "means plus function" claim, as permitted by 35 U.S.C. § 112, sixth paragraph. As such, the Specification must simply provide at least one example of a corresponding structure that performs the means plus function limitations. The Applicants respectfully submit that at least FIGs. 1 and 3 of the

Applicants' Drawings, and the corresponding portions of the Applicants' Specification, clearly describe corresponding structure for the claimed features.

The Examiner submits that the "means" recited in claims 13-14 could, under the broadest reasonable interpretation, be "interpreted as software only" (Office Action, Page 5). Notably, it is irrelevant under 35 U.S.C. § 112, sixth paragraph if a means plus function limitation could be a software apparatus. Rather, as long as the Applicants' Specification provides <u>at least one corresponding structure</u>, the requirements of 35 U.S.C. § 112, sixth paragraph are fulfilled. Thus, the Applicants submit that claims 13-14 fully satisfy the requirements of both 35 U.S.C. § 101 and 35 U.S.C. § 112, sixth paragraph.

Moreover, the Applicants respectfully direct the Examiner's attention to the "Interim Examination Instructions For Evaluating Subject Matter Eligibility Under 35 U.S.C. § 101," effective August 24, 2009. Page 5 of the Interim Examination Instructions states that "[a]n 'apparatus' does not have a significantly different meaning from a machine and can include a machine or group of machines or <u>a totality of means by which a designated function or specific task is executed</u>" (emphasis added). Therefore, the Applicants respectfully submit that the Interim Examination Instructions clearly support subject matter eligibility under 35 U.S.C. § 101 for means plus function claims.

Finally, according to the two-step § 101 analysis set forth in the slides that accompany the Interim Examination Instructions, the first step in determining the subject matter eligibility of a claim under 35 U.S.C. § 101 is to determine whether the claim is "directed to ... [a] process, machine, manufacture or composition of matter" (Slide 4). Claims 13-14 clearly recite "<u>an intrusion detection system</u> that includes <u>a plurality of sensors</u>." As stated above, Page 5 of the Interim Examination Instructions states that "[a]n 'apparatus' (*e.g.*, an intrusion detection system and/or a sensor) does not have a significantly different meaning from a <u>machine</u> ..." (emphasis added). As such, the Applicants submit that claims 13-14 are directed to an apparatus or machine, and, as such, clearly satisfy the first step of the two-step analysis.

The second step of the two-step analysis specifies that "[a] claim satisfying Step 1 is subject matter eligible under 101 unless it wholly embraces a judicially recognized exception," where those judicially recognized exceptions include abstract ideas, laws of

nature, natural phenomena, mental processes, mathematical algorithms, and scientific principles (Slide 5). Since claims 13-14 clearly do not wholly embrace a judicially recognized exception, the Applicants submit that, according to the two-step analysis, claims 13-14 are clearly subject matter eligible under 35 U.S.C. § 101.

In light of the above, the Applicants respectfully submit that claims 13-14 are clearly directed to statutory subject matter and request that the rejection of claims 13-14 under 35 U.S.C. § 101 be withdrawn.

II. REJECTION OF CLAIMS 1-2, 7-8, AND 13-14 UNDER 35 U.S.C. § 102

Claims 1-2, 7-8, and 13-14 stand rejected as being unpatentable over the Nine et al. patent (U.S. 6,560,611, issued May 6, 2003, hereinafter "Nine"). The Applicants respectfully traverse the rejection. Specifically, the Applicants submit that Nine fails to teach, show, or suggest several of the features recited in Applicants' independent claims 1, 7, and 13.

Primarily, the Applicants submit that Nine is completely devoid of any teaching, showing, or suggestion relating to the comparison of an alert (indicating an attack or anomalous incident) – or more specifically, the comparison of features of the alert - to the features of existing alert classes, in order to classify the alert, as claimed by the Applicants in independent claims 1, 7, and 13.

By contrast, Nine teaches a network monitoring system that simply reports a detected problem to the proper individual (e.g., technician), based on the nature of the problem. That is, Nine does not classify the detected problem (e.g., in accordance with its features) by comparing it to known problems, but simply evaluates the detected problem as a discrete incident and reports it to a human technician for further action.

Specifically, Nine teaches a remote monitoring system (RMS) that reports to a network operation site (NOS) when the RMS detects an anomaly with respect to a service it monitors. The report provided by the RMS is a ticket or data record containing information about the service (e.g., location, severity of problem, time of occurrence). In addition, the system "determines the nature of the problem, and notifies the proper personnel [e.g., a technician]" (See, Nine at column 3, lines 25-27).

For instance, the portion of Nine that the Examiner cites to teach the limitations

of "updating a threshold similarity expectation for one or more features" of an alert relative to features of alert classes and of "updating a similarity expectation for one or more features" of an alert relative to features of alert classes in fact merely teaches that the monitoring software is replicated for each service on a device by an informer engine executing forker software and sender software (See, e.g. Nine at column 5, line 45 – column 6, line 9). There is no discussion of examining the features of an alert, or of the need to update a threshold similarity requirement or a similarity expectation for the features of the alert to the one or more alert classes.

The portion of Nine that the Examiner cites to teach the limitations of "comparing [a] new alert with one or more alert classes" and "associating the new alert with the existing alert class that the new alert most closely matches" in fact merely teaches three techniques for detecting a problem with a monitored service. The first technique checks to make sure that the service is responsive (e.g., by "ping, nmap, finger, or telnet", Nine at column 7, lines 25-33). The second technique monitors environmental sensors to detect problems with the environment (e.g., "if the temperature is too high", Nine at column 7, lines 34-39). The third technique examines a log of the monitored service and parses for potential problems (e.g., indication that a particular route associated with a router is not functioning, Nine at column 7, lines 40-46). None of these techniques involve the comparison of an alert to existing alert classes, or the association of the alert with one of the existing alert classes based on the comparison.

The portion of Nine that the Examiner cites to teach the limitations of "comparing [a] new alert with one or more alert classes" and "defining a new alert class that is associated with the new alert" in fact merely teaches that log files for a monitored service may be used to diagnose problems with the service. Again, there is no mention of the need to compare an alert with existing alert classes in order to classify the alert, as claimed by the Applicants.

Moreover, Nine does not even teach, show, or suggest that an alert may be classified in accordance with its features. The portion of Nine that the Examiner cites to teach the limitation of "identifying a set of potentially similar features shared by [a] new alert and one or more existing alert classes" in actuality merely teaches that software monitors a service and reports to the NOS when the service is unresponsive or when an anomaly is detected. The report contains "information about the service, such as

location, severity of the problem, and time of occurrence" (*See, e.g.*, Nine at column 3, lines 12-20). There is no mention in this passage of the need to <u>identify features of the problem</u> or to <u>compare the problem to other known problems</u> (*e.g.*, existing alert classes) based on the identified features.

In short, as discussed above, Nine fails to teach, show, or suggest any sort of classification of alerts by <u>comparing features of the alerts to features of existing alert classes</u>, as recited by the Applicants in independent claims 1, 7, and 13. The Examiner appears to suggest that because Nine parses a received ticket in order to determine where to place it, Nine performs a classification step. Specifically, the Examiner submits that "[i]n parsing the ticket, the receiver is taking features of the alert then comparing them to other alerts and classifying the alert, which is deciding where to place the pending ticket. In other words, the pending ticket gets placed with similar pending tickets, which are those in the same class. The class is decided by comparing the features of the ticket to features of other tickets" (Office Action, Page 3). The Applicants respectfully disagree.

First, as discussed in the Applicants' previous response of November 24, 2008, the Applicants submit that the Examiner is reading far more into the teachings of Nine than is actually disclosed. The portion of Nine that the Examiner cites to support the above argument merely states that "[u]pon receipt of the ticket, receiver process 250 parses the ticket and uses the information in the ticket to query accounting engine 248 for information on where to place the pending ticket" (Nine, column 8, lines 38-41). It requires a significant intuitive leap to suggest that deciding where to place a ticket is the same as classifying the pending ticket by comparing it to other pending tickets. There are many possible ways in which the proper location for a pending ticket could be determined. The above sentence teaches determining where to place a pending ticket based on <u>information contained in the pending ticket</u>; nothing in this sentence even alludes to the possibility that <u>information contained in other tickets</u> may be useful in determining where to place the pending ticket. Thus, there is simply no support in Nine for the step of comparing a pending ticket to other tickets for classification purposes.

Reading further into the paragraph cited by the Examiner, Nine further teaches that the accounting engine that is queried for information on where to place the pending ticket "queries administrator file 262 for information regarding the service" (Nine, column

8, lines 44-46). In particular, the accounting engine "reads administrator file 262 from database 258 for this information [the location to place an incoming ticket]" (Nine, column 7, lines 13-18). Specifically, the accounting engine "queries administrator file 262 and responds with the technical or sales personnel to be notified of the pending ticket and the method of notification to use" (Nine, column 8, lines 50-53, emphasis added). Thus, in determining where to place the pending ticket, Nine simply matches information about the pending ticket (e.g., IP or port address of a nonresponsive web service) to the person who is tasked with resolving the indicated problem. This does not involve or require comparing the pending ticket to other tickets, as claimed by the Applicants.

Moreover, the Applicants respectfully submit that the explicit teachings of Nine actually teach away from the claimed classification step. Specifically, Nine teaches that an additional reporting feature is required to detect groups of tickets related to a common problem (See, e.g., Nine, column 9, lines 30-39: "if a series of tickets indicate that a security log file on an NT server has a flood of ICMP packets, a report may be created to locate all of the tickets that indicate this problem," emphasis added). If the tickets had been classified (i.e., compared against other tickets and grouped together into classes) as they were generated (i.e., before being transmitted to the appropriate location), then such a report would not be necessary. That is, all of the tickets that indicate the problem would already be grouped together. Thus, the post-transmission reporting feature required by Nine clearly indicates that Nine teaches that the tickets are not classified or compared to each other, as claimed by the Applicants.

Specifically, Applicants' claims 1, 7, and 13 positively recite:

1.   In an intrusion detection system that includes a plurality of sensors that generate alerts when attacks or anomalous incidents are detected, a method for organizing the alerts into alert classes, both the alerts and the alert classes having a plurality of features, the method comprising the steps of:
    (a) receiving a new alert;
    (b) identifying a set of similar features shared by the new alert and one or more existing alert classes;
    (c) updating a threshold similarity requirement for one or more of the similar features;
    (d) updating a similarity expectation for one or more of the similar features;
    (e) comparing the new alert with the one or more existing alert classes; and either:

(f1) <u>associating the new alert with a one of the one or more existing alert classes that the new alert most closely matches</u>; or

(f2) <u>defining a new alert class that is associated with the new alert</u>. (Emphasis added)

7.   A computer readable medium containing an executable program for organizing alerts that are generated by a plurality of sensors into alert classes, both the alerts and the alert classes having a plurality of features, where the program performs the steps of:

(a) receiving a new alert;

(b) <u>identifying a set of similar features shared by the new alert and one or more existing alert classes</u>;

(c) <u>updating a threshold similarity requirement for one or more of the similar features</u>;

(d) <u>updating a similarity expectation for one or more of the similar features</u>;

(e) <u>comparing the new alert with the one or more existing alert classes</u>; and either:

(f1) <u>associating the new alert with a one of the one or more existing alert classes that the new alert most closely matches</u>; or

(f2) <u>defining a new alert class that is associated with the new alert</u>. (Emphasis added)

13.   In an intrusion detection system that includes a plurality of sensors that generate alerts when attacks or anomalous incidents are detected, a system for organizing the alerts into alert classes, both the alerts and the alert classes having a plurality of features, where the system comprises:

(a) means for receiving a new alert;

(b) <u>means for identifying a set of similar features shared by the new alert and one or more existing alert classes</u>;

(c) <u>means for updating a threshold similarity requirement for one or more of the similar features</u>;

(d) <u>means for updating a similarity expectation for one or more of the similar features</u>;

(e) <u>means for comparing the new alert with the one or more existing alert classes</u>; and

(f1) <u>means for associating the new alert with a one of the one or more existing alert classes that the new alert most closely matches, or defining a new alert class that is associated with the new alert</u>. (Emphasis added)

As discussed above, nowhere does Nine teach or even suggest the desirability of classifying of alerts by <u>comparing features of the alerts to features of existing alert classes</u>. Moreover, even assuming for the sake of argument that the disclosure of Nine can be interpreted as teaching the classification of alerts, Nine fails to teach or suggest several other claimed features of the present invention, namely, the steps of <u>updating a</u>

threshold similarity expectation for one or more features of an alert relative to features of alert classes and of updating a similarity expectation for one or more features. Therefore, the Applicants submit that independent claims 1, 7, and 13 fully satisfy the requirements of 35 U.S.C. §102 and are patentable thereunder.

Dependent claims 2, 8, and 14 depend, respectively, from claims 1, 7, and 13, and recite additional features therefore. As such, and for at least the same reasons set forth above, the Applicants submit that claims 2, 8, and 14 are not anticipated by the teachings of Nine. Therefore, the Applicants submit that dependent claims 2, 8, and 14 also fully satisfy the requirements of 35 U.S.C. §102 and are patentable thereunder.

III. CONCLUSION

Thus, the Applicants submit that all of the presented claims fully satisfy the requirements of 35 U.S.C. §101 and 35 U.S.C. §102. Consequently, the Applicants believe that all of these claims are presently in condition for allowance. Accordingly, both reconsideration of this application and its swift passage to issue are earnestly solicited.

If, however, the Examiner believes that there are any unresolved issues requiring the issuance of a final action in any of the claims now pending in the application, it is requested that the Examiner telephone Kin-Wah Tong, Esq. at (732) 842-8110 so that appropriate arrangements can be made for resolving such issues as expeditiously as possible.

Respectfully submitted,

December 11, 2009
Date

Kin-Wah Tong, Attorney
Reg. No. 39,400
(732) 842-8110

Wall & Tong, LLP
595 Shrewsbury Avenue
Shrewsbury, New Jersey 07702